

What is claimed is:

- 1           1.     A database system comprising:  
2                     one or more storage devices containing a table having plural rows,  
3                     the plural rows including a first row containing a before image  
4     representing data before a data modification operation and a second row containing an  
5     after image representing data processed by the data modification operation.
- 1           2.     The database system of claim 1, wherein the one or more storage devices  
2     further contain identifiers to identify a state of each row.
- 1           3.     The database system of claim 2, wherein the identifiers are contained in  
2     the table.
- 1           4.     The database system of claim 1, wherein the first and second rows are  
2     associated with the same row identifier.
- 1           5.     The database system of claim 4, wherein the table further contains state  
2     identifiers to identify a before image state of the first row and an after image state of the  
3     second row.
- 1           6.     The database system of claim 1, wherein the data modification operation is  
2     performed in a transaction, the transaction having one or more requests, wherein the first  
3     row contains a transaction before image representing data before the beginning of the  
4     transaction, and wherein the plural rows further comprise a third row containing a request  
5     before image representing data before the beginning of a request in the transaction.
- 1           7.     The database system of claim 5, further comprising a module adapted to  
2     transition the state of each row based on a data manipulation command.

1           8.       The database system of claim 1, further comprising a module adapted to  
2 return data in the second row in response to a read request under a normal condition.

1           9.       The database system of claim 8, wherein the module is adapted to return  
2 data in the first row in response to a read request under an abort condition in which the  
3 modification operation is aborted.

1           10.      The database system of claim 9, further comprising a rollback module  
2 adapted to mark the first row as containing a current image in response to the abort  
3 condition.

1           11.      The database system of claim 10, wherein the rollback module is adapted  
2 to further remove the second row in response to the abort condition.

1           12.      The database system of claim 1, wherein the table contains a first row  
2 identifier associated with the first and second rows, a first state identifier having a first  
3 value associated with the first row, and a second state identifier having a second value  
4 associated with the second row.

1           13.      The database system of claim 12, wherein the table further contains a  
2 mutation identifier associated with the first row identifier to identify that the modification  
3 operation is occurring with respect to one or more rows associated with the first row  
4 identifier.

1           14.      The database system of claim 13, wherein the mutation identifier changes  
2 value with each new modification operation.

1           15.      The database system of claim 14, wherein the data modification operation  
2 is performed in a transaction, each transaction having one or more requests, the mutation  
3 identifier having a transaction identifier portion and a request identifier portion.

1           16.     The database system of claim 15, wherein the transaction identifier portion  
2     has a value that increments with each new transaction.

1           17.     The database system of claim 14, further comprising a module adapted to  
2     return a row to return based on the mutation identifier and state identifier information.

1           18.     The database system of claim 14, wherein the one or more storage devices  
2     further contain an active mutation identifier list having one or more mutation identifiers  
3     associated with one or more active modification operations.

1           19.     The database system of claim 18, wherein the one or more storage devices  
2     further contain an abort mutation identifier list having one or more mutation identifiers  
3     associated with one or more aborts of modification operations.

1           20.     A method of providing access in a database system, comprising:  
2                   storing data in rows of a table; and  
3                   in response to a data modification operation of a first row, marking the  
4     first row as a before image row containing data before the start of the data modification  
5     operation, and creating a second row as an after image containing data processed by the  
6     data modification operation.

1           21.     The method of claim 20, further comprising setting a first state identifier  
2     to a first value to identify the first row as the before image row and setting a second state  
3     identifier to a second value to identify the second row as the after image row.

1           22.     The method of claim 21, further comprising returning the second row in  
2     response to a read operation under a first condition.

1           23.     The method of claim 22, further comprising returning the first row in  
2 response to the read operation under a second condition in which the data modification  
3 operation has been aborted.

1           24.     The method of claim 20, further comprising rolling back to the first row if  
2 the data modification operation aborts.

1           25.     The method of claim 24, further comprising deleting or marking as  
2 available for reuse the second row during a rollback process in response to the abort.

1           26.     The method of claim 20, further comprising marking the second row as a  
2 current row if the data modification operation commits.

1           27.     The method of claim 26, further comprising deleting or marking as  
2 available for reuse the first row once the data modification operation commits.

1           28.     An article comprising at least one storage medium containing instructions  
2 that when executed cause a system to:

3                     store data in rows of a table; and  
4                     store a state identifier associated with each row, the state identifier having  
5 a first value to indicate a row as being a before image of a data modification operation  
6 and a second value to indicate a row as being an after image of a data modification  
7 operation.

1           29.     An article comprising at least one storage medium containing:  
2                     a data structure having plural portions,  
3                     the data structure further containing state identifiers associated with  
4 corresponding portions, a first state identifier having a first value to indicate a row as  
5 being a before image of a data modification operation and a second state identifier having  
6 a second value to indicate a row as being an after image of a data modification operation.